Bank of Canada Banque du Canada

Working Paper 2002-18 / Document de travail 2002-18

Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data

by

Kevin Moran and Veronika Dolar

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-18

July 2002

Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data

by

Kevin Moran and Veronika Dolar

Monetary and Financial Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
kmoran@bank-banque-canada.ca
vdolar@bank-banque-canada.ca

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements iv
Abstract/Résumé v

1. Introduction 1

2. The Model 2

   2.1 The one-sector, one-shock RBC model 2

   2.2 Solving the structural model 4

3. Estimation 5

   3.1 Estimation procedure 5

   3.2 The data 7

4. Results 7

   4.1 A benchmark for comparison: a two-lag, reduced-form VAR 7

   4.2 Full-sample estimation 8

   4.3 Stability of the parameters 10

   4.4 Out-of-sample forecasting accuracy 11

5. Discussion 13

   5.1 Why the hybrid model performs better: within-sample versus out-of-sample fit 13

   5.2 Future research 14

6. Conclusion 15

References 17

Tables 20

Figures 27

Appendix A: Computing the Log-Likelihood of the Model using the Kalman Filter: A Brief Overview 30

Acknowledgements

Thanks to Peter Ireland for making his programs available to us and to Robert Amano, Ali Dib, Céline Gauthier, Scott Hendry, Jinill Kim, and Christian Zimmermann for useful discussions and comments. We also wish to thank Eli Papakirykos for his assistance on a previous version of the paper.

Abstract

This paper applies the hybrid dynamic general-equilibrium, vector autoregressive (DGE-VAR) model developed by Ireland (1999) to Canadian time series. It presents the first Canadian evidence that a hybrid DGE-VAR model may have better out-of-sample forecasting accuracy than a simple, structure-free VAR model. The evidence suggests that estimated DGE models have the potential to add good forecasting ability to their natural strength of adding structure to an economic model.

JEL classification: E32, E37
Bank classification: Business fluctuations and cycles; Economic models

Résumé

Dans cette étude, les auteurs appliquent aux séries chronologiques canadiennes le modèle hybride dynamique d'équilibre général et vectoriel autorégressif conçu par Ireland (1999). Les résultats montrent pour la première fois, dans le cas du Canada, que l'exactitude des prévisions hors échantillon peut être supérieure dans un modèle hybride de ce type que dans un modèle VAR simple sans structure. Les observations laissent croire que les modèles dynamiques d'équilibre général estimés peuvent allier un bon pouvoir de prévision à leur capacité naturelle de structurer un modèle économique.

Classification JEL : E32, E37
Classification de la Banque : Cycles et fluctuations économiques; Modèles économiques

102-06650/

1. Introduction

Ireland (1999) develops a hybrid dynamic general-equilibrium, vector autoregressive (DGE-VAR) model. The model is a hybrid because it supplements a one-sector, one-shock real-business-cycle (RBC) model with three non-structural sources of fluctuations. Ireland presents evidence that the parameters related to the structural part of the model are stable, whereas those originating in its non-structural part are not. Further, he shows that the model's out-of-sample forecasting ability is superior to that of a simple, reduced-form VAR.

In this paper, we apply Ireland's analysis to Canadian data. We replicate his results on the stability of the structural parameters and we report the first Canadian evidence that a hybrid model can outperform a simple, reduced-form VAR in out-of-sample forecasting. Our results thus support the view that macroeconomic models built on DGE foundations, along with their natural strength in interpreting macroeconomic data in terms of stable, structural features—preferences, technology, and optimization—have the potential to deliver good forecasting ability.

The strategy of estimating a DGE-based model by maximum likelihood, rather than by calibrating it, agrees with a rapidly developing literature that chooses to anchor the quantitative expression of DGE models within standard econometrics.[1] We believe that using such techniques to confront the model with the data improves the identification of the structural parameters, facilitates comparisons with alternative models, and extends the reach of the DGE methodology.

Ireland's (1999) paper and our paper contribute to this literature by assessing the out-of-sample forecasting ability of models based on DGE methodology. This ability has been examined before, within the Bayesian frameworks of Ingram and Whiteman (1994) and DeJong, Ingram, and Whiteman (2000), and found to be surprisingly good, relative to popular alternatives.[2] Such results are similar to ours, which we obtain using a very different econometric strategy. Together, the results suggest that good forecasting performance might be a robust feature of DGE-based models.

By relying partly on non-structural sources of fluctuations to capture the dynamics of the time series examined, we chose not to attempt a complete structural characterization of macroeconomic data. This does not mean that such attempts should not be undertaken. Rather, it recognizes that central banks need good forecasting models for policy analysis in the present as well as continuing research to build the models of the future. Our strategy suggests that a model where structural and non-structural elements co-exist might constitute, at a given time, the best

[1] See Leeper and Sims (1994), McGrattan, Rogerson, and Wright (1997), Ireland (1997, 2001a,b), Kim (2000), and Dib (2001, 2002).

[2] In a related paper, Altig, Carlstrom, and Lansing (1995) compare the out-of-sample fit of a calibrated DGE model with those of two alternatives: a reduced-form VAR and the White Book projections.

policy model.

This paper is organized as follows. Section 2 reviews the one-sector, one-shock RBC model and its first-order, approximate solution. Section 3 describes the procedure used to estimate the model. This procedure, suggested by Sargent (1989), uses the Kalman filter to recursively construct the model's likelihood. It is at this stage that non-structural sources of fluctuations are added to the original RBC structure. Section 4 describes the results on estimation, on the stability of the parameters, and on the out-of-sample forecasting accuracy of the model. Section 5 discusses the results. We conjecture that the parsimony inherent in the structural model may be related to its good relative forecasting accuracy. We also propose different dimensions along which future work could broaden the results described here.

2. The Model

2.1 The one-sector, one-shock RBC model

Households seek to maximize expected, discounted lifetime utility, subject to a sequence of budget constraints.[3] The optimization problem is the following:

$$\max_{\{c_{t+i},n_{t+i},k_{t+1+i}\}_{i=0}^{\infty}} E_t \sum_{i=0}^{\infty} \beta^i (ln(c_{t+i}) + \psi(1 - n_{t+i})), \tag{1}$$

subject to:

$$c_t + k_{t+1} - (1-\delta)k_t \leq r_t k_t + w_t n_t, \ \forall t, \tag{2}$$

where c_t and n_t denote consumption and hours worked, respectively, while r_t and w_t denote the rental rates on capital and labour. In addition, the discount factor satisfies $0 < \beta < 1$; the weight on leisure, ψ, is greater than 0; the depreciation rate satisfies $0 < \delta < 1$; and the Lagrangian multiplier of the budget constraint is λ_t.

The first-order conditions with respect to c_t, n_t, and k_{t+1} that result from this optimization problem are as follows:

$$1/c_t = \lambda_t; \tag{3}$$

$$\psi = \lambda_t w_t; \tag{4}$$

$$\lambda_t = \beta E_t [\lambda_{t+1}(r_{t+1} + 1 - \delta)]. \tag{5}$$

[3]The exposition of the standard RBC model presented here is brief. Moran (2002) provides additional details about the model's derivation and the first-order, approximate solution method used most often to characterize its equilibrium.

Firms operate the economy's technology and produce final output, y_t. They are constrained by the familiar constant-return-to-scale (CRS) production function:

$$y_t = F(k_t, n_t) = A_t k_t^{\alpha} (\eta^t n_t)^{1-\alpha}, \quad 0 < \alpha < 1, \tag{6}$$

where η represents deterministic technological improvements and A_t expresses transitory technology shocks around that trend.[4] A_t evolves according to the following AR(1) process:

$$ln(A_t) = (1-\rho)ln(A) + \rho ln(A_{t-1}) + \xi_t, \;\; \xi_t \sim N(0, \sigma_\xi), \tag{7}$$

where A denotes the long-run-mean, $\rho(< 1)$ determines the persistence, and ξ_t denotes a (serially uncorrelated) innovation normally distributed with mean zero and standard deviation σ_ξ.

Firms equate the marginal products of labour and capital to their respective rental rates, so that:

$$(1-\alpha) A_t k_t^{\alpha} \eta^{t(1-\alpha)} n_t^{-\alpha} = w_t; \tag{8}$$

$$\alpha A_t k_t^{\alpha-1} (\eta^t n_t)^{1-\alpha} = r_t. \tag{9}$$

Using these two equations to eliminate w_t and r_t from the households' first-order conditions and from the budget constraint, we can describe the equilibrium of this model with the following five equations and five unknowns, y_t, c_t, n_t, k_t, and λ_t:

$$1/c_t = \lambda_t; \tag{10}$$

$$\psi = \lambda_t (1-\alpha) A_t k_t^{\alpha} \eta^{t(1-\alpha)} n_t^{-\alpha}; \tag{11}$$

$$\lambda_t = \beta E_t \left[\lambda_{t+1} (\alpha A_{t+1} k_{t+1}^{\alpha-1} (\eta^{t+1} n_{t+1})^{1-\alpha} + 1 - \delta) \right]; \tag{12}$$

$$c_t + k_{t+1} - (1-\delta) k_t = y_t; \tag{13}$$

$$y_t = A_t k_t^{\alpha} (\eta^t n_t)^{1-\alpha}. \tag{14}$$

These equations depend on the values of eight structural parameters describing preferences and technology: $\beta, \delta, \psi, \alpha, \eta, A, \rho$, and σ_ξ.

[4] The separation of technology into a deterministic trend, η, and a stationary shock, A_t, implies that output, consumption, and investment are also stationary around a linear trend. This assumption contrasts with results from the empirical literature, suggesting that many economic time series can be well characterized by integrated processes. See section 5 for an additional discussion.

2.2 Solving the structural model

For given values of the eight parameters, the model can be solved using a first-order approximation around the non-stochastic steady-state values. The first step in such a solution is thus to compute steady-state values of the five variables. This is done by eliminating all time subscripts from equations (10) to (14), fixing A_t to its long-run mean, A, and solving the five equations in five unknowns.[5]

Next, a first-order Taylor approximation of the five equations (plus equation (7), which describes the evolution of the productivity shock, A_t) around the steady state is computed. Each equation is thus rewritten in a linearized form, with variables now expressed as deviations from their steady-state values. For example, the linearized version of (10) is:

$$\widehat{\lambda}_t + \widehat{c}_t = 0, \tag{15}$$

where $\widehat{\ }$ expresses the percentage deviation of the variable from its steady-state value.[6]

Once the six equations are linearized, algorithms presented in Blanchard and Kahn (1980), King, Plosser, and Rebelo (1988a), or King and Watson (1998) are used to transform the forward-looking system into one expressing all variables as a linear function of state (predetermined) variables and of the innovation to the exogenous shocks. In general terms, the solution has the following form:

$$\begin{aligned} \mathbf{s}_{t+1} &= \mathbf{M_1 s}_t + \mathbf{M_2}\epsilon_{t+1}; \\ \mathbf{f}_t &= \mathbf{\Pi s}_t, \end{aligned} \tag{16}$$

where $\mathbf{s}_t$ represents the state variables of the system, ϵ_{t+1} is the innovation to the exogenous shocks, and $\mathbf{f}_t$ represents all other endogenous (flow) variables. In our example, we would have $\mathbf{s}_t = (\hat{k}_t, \hat{A}_t)$, $\epsilon_{t+1} = \xi_{t+1}$, and $\mathbf{f}_t = (\hat{y}_t, \hat{c}_t, \hat{n}_t, \hat{\lambda}_t)'$.

The matrices $\mathbf{M_1}, \mathbf{M_2}$, and $\mathbf{\Pi}$ are known (non-linear) functions of the structural parameters of the economy. That is, once values for these parameters have been established, the matrices and the system in (16) represent a complete solution of the model. One can then proceed to any simulation desired, tracing out, for example, impulse responses following a 1 per cent shock to ξ_t.

A calibration strategy for the model would thus proceed by assigning values to these parameters. Such values could be taken from other studies, chosen so that the steady-state version of the model replicates long-term averages in the data, or even be arbitrary, if one were interested in exploring the sensitivity of the solution to those values.

---

[5] In addition, the four variables that exhibit growth in the steady state (y_t, c_t, k_t, and λ_t) must first be deflated to ensure the existence of such a steady state. Note that the assumption of trend stationarity makes the detrending straightforward. See Moran (2002) for further details.

[6] We thus have $\widehat{c}_t = \frac{(c_t - c_{ss})}{c_{ss}}$. Ireland uses a slightly different notation, in which the deviation of consumption from its steady-state value would be noted as $\widehat{c}_t = ln(c_t/c_{ss})$.

3. Estimation

3.1 Estimation procedure

The estimation of the model seeks to assign values to the parameters by confronting the solution in (16) with the data. Anticipating the fact that the estimation will be conducted using a state-space framework, the first step in the estimation procedure is to interpret the two equations of (16). The first equation (now labelled the "state" equation) describes the evolution of unobserved state variables of the economy. The second equation (now labelled the "observation" equation) describes how the observed flow variables depend only on the values of the state variables.

The second step is to realize that not all variables in the vector $\mathbf{f}_t$ may actually be observed. In our case, while we do have data for output, consumption, and hours, we obviously do not have data on $\hat{\lambda}_t$, the Lagrangian multiplier. We need to redefine the vector $\mathbf{f}_t$ so that it contains only variables for which we have data. This, in turn, requires that we drop the lines of the matrices $\mathbf{\Pi}$ that correspond with the dropped variable(s) in $\mathbf{f}_t$. Our redefined vector now reads $\mathbf{f}_t^* = (\hat{y}_t, \hat{c}_t, \hat{n}_t)'$, while the redefined matrix $\mathbf{\Pi}^*$ consists of the first three lines of $\mathbf{\Pi}$.[7]

If one of the structural parameters appears only in the line of $\mathbf{\Pi}$ that is dropped, the econometric model will not be able to estimate that parameter. Even if the deleted line of $\mathbf{\Pi}$ does not preclude the estimation of a parameter, it might seriously diminish the efficiency of the estimation.

To this point the solution exhibits only one type of shock, (ξ_t) for three observed variables ($\hat{y}_t, \hat{c}_t$, and $\hat{n}_t$), and as such the econometric model contains linear combinations of the variables that hold exactly. Confronted with data where these combinations do not hold, an estimation procedure based on maximum likelihood would break down.

To circumvent this singularity issue, consider augmenting the model in (16) by adding serially correlated residuals, as in the following:

$$\begin{aligned} \mathbf{s}_{t+1} &= \mathbf{M_1 s_t} + \mathbf{M_2}\epsilon_{t+1}; \\ \mathbf{f}_t^* &= \mathbf{\Pi}^* \mathbf{s}_t + \mathbf{u}_t; \\ \mathbf{u}_{t+1} &= \mathbf{D u_t} + \nu_{t+1}, \quad \nu_{t+1} \sim \mathbf{N(0, V)}. \end{aligned} \tag{17}$$

Alternative interpretations can be given to the residuals $\mathbf{u}_t$. Sargent (1989) and McGrattan, Rogerson, and Wright (1997) view them as measurement errors in the endogenous variables $\mathbf{f}_t$. Such an interpretation naturally leads to the assumption

[7] We might also have data on some of the state variables, although they are assumed to be unobservable by the state-space framework. We could have, for example, constructed a quarterly series for physical capital (k_t). Such a situation could be accommodated by adding a variable to the vector $\mathbf{f}_t$ and another line to $\mathbf{\Pi}$. We thus would have $\mathbf{f}_t^* = (\hat{y}_t, \hat{c}_t, \hat{n}_t, \hat{k}_t)'$ and the vector $[1, 0]$ added to the bottom of the matrix $\mathbf{\Pi}$, so that the equality in the observation equation is preserved.

that the matrices $\mathbf{D}$ and $\mathbf{V}$ are diagonal. It also shapes the view that these errors are not an interesting (economically) source of the dynamics present in the data.

Conversely, Ireland (1999) interprets the errors $\mathbf{u}_t$ as capturing all data dynamics not explained by the structural model. He thus views $\mathbf{u}_t$ not as a noise process merely enabling the estimation to proceed, but as a *bona fide* part of the model, contributing substantially to its stochastic structure.[8]

One interpretation of Ireland's view could be that the residuals $\mathbf{u}_t$ are a stand-in for shocks arising from the demand side of the economy (monetary policy and fiscal shocks), and not modelled within the RBC framework. Alternatively, they could be interpreted as imperfections in the optimizing process underlying solution (16), which decay at rates governed by the matrix $\mathbf{D}$.[9]

Ireland's modelling strategy shares some similarities with the procedure used to build the Quarterly Projection Model (QPM), in which an optimizing core structure is supplemented by estimated dynamic elements, enabling the model to better match data.[10] Ireland's approach is distinguished by the separation of the structural and the non-structural parts of the model and the simultaneous and systematic estimation of the complete model using standard econometric procedures.

To estimate (17), rewrite it as follows:

$$\begin{aligned} \begin{bmatrix} \mathbf{s}_{t+1} \\ \mathbf{u}_{t+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{M}_1 & 0 \\ 0 & \mathbf{D} \end{bmatrix} \begin{bmatrix} \mathbf{s}_t \\ \mathbf{u}_t \end{bmatrix} + \begin{bmatrix} \mathbf{M}_2 & 0 \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \epsilon_{t+1} \\ \nu_{t+1} \end{bmatrix} \\ \mathbf{f}_t^* &= \begin{bmatrix} \mathbf{\Pi}^* & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{s}_t \\ \mathbf{u}_t \end{bmatrix}. \end{aligned} \tag{18}$$

This structure has the form of a state-space representation (see Hamilton 1994, chapter 13), with the state equation on the top and the observation equation on the bottom (recall that the vector $\mathbf{f}_t^*$ now contains only $\hat{y}_t, \hat{c}_t$, and $\hat{n}_t$). From this representation, the Kalman filter can be used to evaluate the log-likelihood for any observed sample $\{\mathbf{f}_t\}_{t=1}^T$.[11] The following estimation strategy is used. First, establish a starting guess for the 23 parameters (8 structural and 15 non-structural, of which 9 are in matrix $\mathbf{D}$ and 6 are in matrix $\mathbf{V}$).[12] Second, use Hamilton's method to

[8] The lag length associated with the data generating process (DGP) for the $\mathbf{u}_t$ variables could be chosen, in principle, according to a likelihood ratio test, rather than being imposed as 1.

[9] Another estimation strategy that circumvents the singularity problem consists of introducing additional structural shocks in the model. Such shocks augment the dimension of the vector ϵ_t until it includes as many structural shocks as there are observed variables. Examples of this estimation strategy and its implications are given in Ireland (1997), Kim (2000), and Dib (2001, 2002).

[10] The QPM is the Bank's main policy model. Black et al. (1994) and Coletti et al. (1996) describe the model in detail.

[11] Appendix A briefly reviews the method by which the filter is used to evaluate the likelihood.

[12] The matrices $\mathbf{D}$ and $\mathbf{V}$ can be written as:

$$\mathbf{D} = \begin{bmatrix} d_{yy} & d_{yc} & d_{yh} \\ d_{cy} & d_{cc} & d_{ch} \\ d_{hy} & d_{hc} & d_{hh} \end{bmatrix} \text{ and } \mathbf{V} = \begin{bmatrix} v_y^2 & v_{yc} & v_{yh} \\ v_{yc} & v_c^2 & v_{ch} \\ v_{yh} & v_{ch} & v_h^2 \end{bmatrix}.$$

evaluate the value of the log-likelihood at those guesses. Third, update each guess in a likely direction of increase in the log-likelihood.[13]

As in Ireland (1999), preliminary estimations report unreasonably low values for β and high values for δ. To retain a plausible structural interpretation, these parameters should, respectively, be close to 1 and close to 0. We thus impose values (standard in the calibration literature) of 0.99 for β and 0.025 for δ, before proceeding to estimate the remaining 21 parameters.[14]

3.2 The data

All the data used for the empirical estimation of the model are taken from Statistics Canada's National Accounts. The series are quarterly and run from 1964Q1 through 2001Q1.

Consumption, C_t, is defined as real personal expenditures on consumer goods and services. Investment, I_t, is defined as real investment (business investment in machinery, equipment, non-residential construction, and residential construction). Output, Y_t, is defined as the sum of C_t and I_t.[15] A series for hours worked, H_t, is obtained by multiplying average hours per week by total employment (summing the monthly series for average hours to a quarterly frequency). Consumption, investment, output, and hours worked are then converted into per-capita terms, using a population of age 15 and over.

4. Results

4.1 A benchmark for comparison: a two-lag, reduced-form VAR

Before reporting on the estimation results of our hybrid model, it is useful to introduce a reduced-form model as a benchmark. The benchmark will help us to assess the within-sample and out-of-sample performances of our model.

[13] We use two types of minimization routine to suggest such directions: the *simplex*, which does not use derivatives, and the quasi-Newton, which does. We repeat all of our numerical experiments with different starting values to identify global optima. Numerical computations are conducted using *Matlab*.

[14] This weakness of our results is not always encountered by researchers estimating DGE models. Both Ireland (2001a) and Kim (2000) report estimates of β comprised between 0.99 and 1. The presence of these calibrated, rather than estimated, parameters reinforces the sense in which the hybrid model considered in this paper is a parsimonious, restricted model, relative to benchmark reduced-form VARs. See section 5 for further discussion.

[15] Since the model abstracts from government activities and expenditures, such a definition of output maximizes the correspondence between the model and data aggregates.

The benchmark is a two-lag VAR in the log-levels of output, consumption, and hours. A constant and a trend are added to the estimation of each equation: the reduced-form model thus embodies the same hypothesis of trend stationarity that is at the base of the hybrid model. The equation for output, for example, has the following form:

$$y_t = \alpha + \beta t + \delta_{11} y_{t-1} + \delta_{12} y_{t-2} + \delta_{21} c_{t-1} + \delta_{22} c_{t-2} + \delta_{31} h_{t-1} + \delta_{32} h_{t-2} + e_t. \quad (19)$$

The lag-length is chosen by likelihood ratio tests that point to the data preferring the two-lag specification over alternatives with one, three, or four lags. Ljung-Box tests for serial correlation of up to four quarters in the residuals indicate that such correlation, present when only one lag is used, disappears once the two-lag specification is adopted. Augmented Dickey-Fuller (ADF) tests indicate that the residuals for the three equations are I(0). The likelihood attained by this model is 1519.47.

A difference-stationary specification of the reduced-form VAR (with or without cointegration) would have been a natural alternative to explore. We do not explore it, to confront the hybrid model with a reduced-form benchmark that rests on the same hypothesis of trend stationarity. In future work, we plan to explore the consequences of assuming a difference-stationary representation, by imposing it simultaneously on the hybrid and the reduced-form model. See section 5.2 for a further discussion.

4.2 Full-sample estimation

Let us now examine the results from the estimation of the hybrid model. Table 1 reports our estimates of the six structural parameters ($\psi, \alpha, \mu, A, \rho$, and σ_ξ) and of the 15 non-structural ones. It also reports standard errors for those estimates as well as their t-statistics.

The structural estimates have economically meaningful values and are highly significant at conventional confidence levels. This high significance mirrors that reported in Ireland (1999). To some extent, it also appears in Dib (2001, 2002), using Canadian data. The estimate $\eta = 1.0038$ implies an annualized, steady-state growth rate of real, per-capita output of 1.5 per cent. The capital share (α) is estimated to be close to 30 per cent, a value similar to those used in the calibration literature. Our estimates of the persistence of technology shocks ($\rho = 0.9860$) and the standard deviation of their innovations ($\sigma_\xi = 0.0054$) are close to those reported by Ireland and those used in the calibration literature. They differ slightly, however, from those reported by Dib (2001, 2002).

The second panel of Table 1, which reports the estimates of the non-structural parameters, illustrates the importance of that portion of the model. In particular, the estimates imply that matrix **D** has two complex eigenvalues of modulus

0.9208 and one, real-valued, equal to 0.5571. These high eigenvalues are the non-structural counterparts to the high persistence of the technology shock. In addition, two standard-deviation estimates (those associated with output, ν_y, and consumption, ν_c) exceed the estimated standard deviation of the innovation to technology. The non-structural fluctuations are thus volatile and persistent, contributing significantly to the model's overall dynamics.[16] Again, on that dimension, our results concur with those reflecting the U.S. experience reported by Ireland.

The likelihood attained by this vector of estimates is 1489.66, which is noticeably lower than the value reported above for the reduced-form model. Although a formal test is not available, the data seem to prefer the reduced-form to the hybrid model.[17]

Table 2 reports the fraction of the k-steps-ahead forecast error variance in output, consumption, investment, and hours worked that results from the technology shocks featured in the model. Panel A of the table indicates that between a third and a half of the forecast variance of output (around the linear trend identified by our estimation) results from the technology shocks. Such levels of explanatory power are lower than those commonly emphasized in the calibration and estimation literature.[18] These lower estimates suggest that domestic technology shocks are less important in explaining fluctuations in the Canadian economy; this finding is consistent with the great openness of the Canadian economy. Dib (2001), however, finds results that agree better with Ireland's. Panel A also shows that the forecasting horizon at which the explanatory power of the technology shocks is lowest, from 8 to 12 quarters ahead, are those most commonly associated with business cycles. Again, Ireland reports similar findings for the U.S. experience.[19]

Table 2 also reports that, as the forecast horizon increases, technology shocks account for an increasing fraction of the forecasting variance of consumption, with this fraction peaking at 80 per cent for the unconditional variance. By contrast, the table also suggests that shorter forecasting horizons of investment and employment are linked to the technology shocks. This feature is particularly striking in the case of employment.

Another way to judge the results of the full-sample estimation consists of visually assessing the fit of the model. Figure 1 shows the in-sample residuals of the model (labelled "hybrid") and compares them with those arising from our benchmark (labelled "reduced-form"). The most striking feature of Figure 1 is that our model does not exhibit wildly different residuals from that of the benchmark, the

[16]The ranking of the estimates ν_y, ν_c, and ν_n matches the ranking in volatility of the deviations from linear trends of the actual data. The estimate ν_n is not, however, statistically different from 0.

[17]Likelihood ratio tests cannot be performed because the two models are not nested.

[18]Prescott (1986) suggests that over 70 per cent of the output variance is explained by technology shocks. Ireland reports estimates suggesting that this figure is around 85 per cent.

[19]This result echoes Watson (1993), in which a spectral decomposition of the output series arising from a simple RBC model reveals the dominance of very low and very high frequencies.

difference in likelihoods notwithstanding.

4.3 Stability of the parameters

To test for the stability of the parameters, two subsamples of equal lengths are employed. The first subsample runs from 1964Q1 through 1982Q3 and the second runs from 1982Q4 through 2001Q1. This break point marks the midpoint of our complete sample.

Table 3 reports estimates of the parameters (along with their standard errors) for the two subsamples. The estimates for ψ, α, and A are almost identical. The estimate of η, however, is higher in the first subsample than in the second (1.0042 rather than 1.0029). This decrease could be interpreted as a reflection of the productivity slowdown. Similarly, the first-sample estimates of ρ and σ_ξ lie slightly above the full-sample estimates; this situation is reversed when the second subsample is used. This result suggests that technology shocks were slightly more volatile and persistent before 1982. Overall, however, the changes in the estimated values of the structural parameters from one subsample to the next appear modest. By contrast, the bottom panel of Table 3 indicates that the estimates for the non-structural parameters exhibit stronger variation across the two samples. This observation suggests that the structural part of the model is relatively invariant across samples, compared with the non-structural portion.

To assess this suggestion, Table 4 reports the results of stability tests. We first test for the stability of all 21 parameters, then for the stability of the 6 structural ones, and then for the stability of the remaining 15 non-structural parameters.[20] Table 4 indicates that, first, we reject the null hypothesis of stability of all 21 estimated parameters at the 1 per cent significance level. Second, we fail to reject the null hypothesis of stability in the 6 structural parameters, at conventional significance levels. Finally, the third test rejects the null of stability of the 15 non-structural parameters. At a minimum, these results suggest that there is considerably less evidence against the stability of the structural parameters than there is against the stability of the non-structural parameters.[21]

[20] The first test is based on a likelihood-ratio argument. The statistic is $LR = 2[lnL(\theta^1) + lnL(\theta^2) - lnL(\theta)]$ with θ^1, θ^2, and θ the full vectors of estimates for the first subsample, the second subsample, and the full sample, respectively. The statistic is asymptotically distributed as a $\chi^2_{(21)}$. The two other tests are Wald tests where the statistic is $W = (\theta^1_q - \theta^2_q)'[cov(\theta^1_q) + cov(\theta^2)_q]^{-1}(\theta^1_q - \theta^2_q)$, where θ^1_q and θ^2_q are estimates of the subset of parameters under study, determined using the first and second subsamples, respectively.

[21] We experimented with a different date for separating the two subsamples: 1979Q4. It leaves the subsamples unbalanced in length but is appealing because it ends a decade of generally rising inflation. While the evidence arising from the experiment was not as clear-cut as the experiment presented in Table 4, the result that there is less evidence against the stability of the structural parameters remained.

4.4 Out-of-sample forecasting accuracy

We compute forecasts of the model over the period from 1985Q1 to 2001Q1. More precisely, we first use data from 1964Q1 through 1984Q4 to estimate the model. Once estimated, the model is used to produce forecasts 1 to 4 quarters ahead; i.e., forecasts for 1985Q1 to 1985Q4. Next, we add the 1985Q1 data point to the original data set, re-estimate the model, and forecast again 1 to 4 quarters ahead; i.e., for 1985Q2 to 1986Q1. We update our estimation and forecasts in such increments until the next-to-last data point (2000Q4) has been added to the estimation. We thus obtain series for one-quarter-ahead forecasts (running from 1984Q4 through 2000Q4), for two-quarter-ahead forecasts (from 1984Q4 through 2000Q3), and so on. These forecasts can then be confronted with observed data.

We also report on the predictive capability of the benchmark VAR introduced in section 4.1. Recall that the benchmark VAR contains no structural identifications of shocks, technology, or preferences. Tables 5 and 6 show the results. Each table reports the forecasting performance—as measured by the root-mean-square of the forecasting errors (RMSE)—for each model and variable, for forecasting horizons of 1 to 4 quarters. Table 5 relates to the full period for which we compare forecasts with realized values (1985Q1 to 2001Q1), whereas Table 6 focuses on the 1990s, to allow a separate analysis of that period.

A glance at the two tables suggests that the RMSEs for the hybrid model are lower than those for the reduced-form benchmark. To determine whether any of the differences are statistically significant, we use the Diebold and Mariano (1995) test on the forecasts from the two models.[22] When the test suggests that one model's RMSE is lower, Tables 5 and 6 signal it with a "D" superscript over the value of the RMSE for that model.

Table 5 (forecasting the period 1985–2001) reports that the hybrid model outperforms the reduced-form model 13 times out of 16, while in the second experiment (Table 6, 1990–2001) it does so 11 times out of 16. The reduced-form model does not outperform the hybrid in any of the 32 comparisons. The results thus suggest a potential advantage in forecasting these series using the hybrid model. Note that in the second experiment (covering the period 1990–2001), the RMSE of the reduced-form model is sometimes lower than that of the hybrid model (in forecasting consumption). These differences are not sufficient, however, to give the reduced-form model an advantage: the Diebold-Mariano statistic does not view this difference as statistically significant. The somewhat weaker results in favour of the hybrid model over the 1990s might be the result of the smaller sample size (we examine the forecasting accuracy over 45, rather than 65, quarters). Alternatively, they could arise because the 1990s, having a more stable monetary policy environment, are easier to forecast

[22] The Diebold and Mariano test assesses the null hypothesis of no difference between the RMSEs of two competing models. It produces a statistic that, under the null, is asymptotically normally distributed.

using a reduced-form VAR.

Tables 5 and 6 report evidence that the hybrid model often has lower RMSEs than the reduced-form model. One still wonders whether forecasts from the reduced-form model contain *any* useful information about future realizations of the variables, over and above that already contained in the forecasts arising from the hybrid model. The forecast-encompassing test, originally devised by Chong and Hendry (1986), is used to answer this question. In effect, it measures whether a combination of the forecasts from the hybrid model and the forecasts from the reduced-form model would improve on the forecasts from the hybrid model alone. When such is *not* the case, we say that the hybrid model *forecast encompasses* the reduced-form model and, in Tables 5 and 6, indicate that fact with a superscript "F" over the RMSE of the hybrid model.[23] Table 5 indicates that in 11 cases out of 16, the hybrid model forecast encompasses the reduced-form model. In Table 6, this happens in 13 cases out of 16. The evidence from this test reinforces the evidence obtained using the Diebold-Mariano test and strongly suggests that the hybrid model posseses a better out-of-sample forecasting ability than the reduced-form model.

The results described in Tables 5 and 6 are expressed visually in Figures 2 and 3, which graph the out-of-sample forecasting error of the hybrid model (labelled "hybrid") and the benchmark ("reduced-form"). Figure 2 shows the errors arising from forecasting one quarter ahead while Figure 3 shows those errors arising from the four-quarters-ahead forecasts.

The first panel in Figures 2 and 3 illustrates the case of output forecasting. Both figures show that, for long periods in the late 1980s, the mid-1990s, and the end of the 1990s, the hybrid model forecasts noticeably better than the benchmark. In the cases of consumption and employment, the improvements of the hybrid over the benchmark are not clearly visible when forecasting one quarter ahead (Figure 2), but become noticeable in the four-quarters-ahead forecasts (Figure 3). Again, the late 1980s and the mid-1990s stand out as periods where the hybrid model performs better than the benchmark.

[23] The test is implemented by the following regression:

$$y_t - \widehat{y}_t^u = \alpha + \beta(\widehat{y}_t^* - \widehat{y}_t^u) + e_t,$$

where $\widehat{y}_t^u$ and $\widehat{y}_t^m$ represent the forecasts from the reduced-form and hybrid model, respectively. The two following null hypotheses are tested: $H_0^1 : \beta = 0$, and $H_0^2 : \beta = 1$. Note the consequence of either null hypothesis being true: the forecasts from one model only are sufficient to explain realized values. If the data reject H_0^1 but not H_0^2, the hybrid model is said to forecast encompass the reduced-form model. Such a result expresses the idea that the combination of the forecasts from both models does *not* perform better than those from the hybrid model alone, but *does* perform better than those from the reduced-form model alone. Conversely, if it were to occur (but it does not) that H_0^2 is rejected but H_0^1 is not, the reduced-form model would forecast encompass the hybrid. The cases (Reject, Reject) and (Accept, Accept) are taken to be inconclusive. This method of implementing the forecast-encompassing test follows Chan and Lafrance (2001), and is slightly different from the original test proposed by Chong and Hendry (1986).

5. Discussion

5.1 Why the hybrid model performs better: within-sample versus out-of-sample fit

Researchers have identified several dimensions along which the empirical implications of simple RBC models are at odds with observed features of the data. For example, Cogley and Nason (1995) show that the simple RBC model, because of its weak internal propagation mechanism, cannot match the autocorrelation function of output or the impulse responses of Blanchard and Quah (1989). Chow and Kwan (1998) demonstrate that the same model, once translated into a VAR in employment, investment, and productivity, implies restrictions on that VAR that are strongly rejected by the data.

Models that extend the simple RBC structure by introducing nominal rigidities and multiple sources of volatility, such as those in Ireland (1997) and Kim (2000), are less at odds with the data. Nevertheless, both researchers report that the likelihood attained by their structural models is lower than that of corresponding reduced-form VARs.[24]

The hybrid model that we consider adds three non-structural shocks to the simple RBC structure, to overcome the weak propagation mechanism singled out by Cogley and Nason (1995). It is therefore expected to possess a good fit with the data. As indicated above, however, the likelihood of the reduced-form model (at 1519.67) is still noticeably higher than the one attained by the hybrid model (which stands at 1489.66).

In such a context, the evidence that structural models may display better out-of-sample forecasting ability, reported in Ireland (1999), our paper, and elsewhere, may seem surprising.[25] Taken generally, this evidence would suggest that restricted or parsimonious specifications (like the hybrid model), although rejected within-sample, may often outperform unrestricted alternatives (the reduced-form model) in out-of-sample exercises. Clements and Hendry (1998, 1999) assess the validity of this conjecture. The main trade-off discussed is that of sampling variability (introduced in the unrestricted specification by the estimation of numerous parameters) versus inconsistency (introduced in parsimonious models by imposing possibly false restrictions). Clements and Hendry conclude that, without frequent structural breaks, parsimony is unlikely to significantly improve forecasting ability. Conversely, the presence of frequent structural breaks leaves open the potential for significant

[24] Recall that no formal likelihood tests can be performed because, in both cases, the structural and the reduced-form models are not nested.

[25] As stated in the introduction, Ingram and Whiteman (1994) and DeJong, Ingram, and Whiteman (2000) display such evidence. In an earlier paper, Ireland (1995) reports that, once translated into a bivariate VAR, the simple version of the permanent income theory is rejected within-sample but helps the model to better forecast out-of-sample.

improvements by imposing some restrictions (among them over-differencing) and better estimating the deterministic elements of the model.

There is no doubt that our dataset contains several instances of structural breaks. The results of section 4.2, where the non-structural parameters are found to have changed significantly across two subsamples, are only one way to characterize this evidence. Therefore, according to Clements and Hendry, our context may be one where a restricted model, while rejected by the data within-sample, can nevertheless outperform the benchmark in out-of-sample forecasting exercises. Furthermore, the estimation of the deterministic trends in the hybrid model proceeds in a manner consistent with Clements and Hendry's prescription: the hybrid model restricts the trend in output and consumption to be the same and the trend in hours to be zero, while the reduced-form benchmark estimates three separate trends for these variables.

The evidence in favour of parsimony is reinforced by conducting the following experiment. Consider using a one-lag reduced-form VAR as the benchmark, rather than the two-lag one used until now. Recall that such a specification is found to be lacking by likelihood ratio tests and by evidence of serial correlation in the residuals of the three equations. Table 7 presents the results of using this VAR as the benchmark. First, a comparison of the lines labelled RMSE (reduced-form model) in Tables 5 and 7 shows that, for short horizons, the RMSEs of the one-lag VAR are larger than those of the two-lag model. As the forecasting horizon increases, however, the instances of where the parsimonious model, the one-lag model, performs better increase. In forecasting relatively long horizons ahead, the model that is found to be lacking using within-sample measurements may be performing better. Further, observe in Table 7 that while the hybrid model still dominates the reduced-form model, the evidence is slightly weaker than it was in Table 5: the Diebold-Mariano test indicates a stronger performance for the hybrid model in 12 cases out of 16, rather than 13 cases out of 16, the result in Table 5.

5.2 Future research

Development of the best possible economic model is a never-ending task. A central bank cannot wait for the perfect model to be constructed to conduct monetary policy. It needs good forecasting models for policy analysis in the present as well. In this context, the set-up in (17) has considerable potential as a bridge between the theoretical and empirical traditions that cohabit within most central banks. Future work could develop as follows.

First, expanding the structural model to include nominal variables, of obvious interest to central banks, is not only possible but by now relatively standard. Ireland (2001a, b) and Dib (2001, 2002), for example, include inflation and money in the vector of flow variables, $\mathbf{f}_t$. In their models, the shocks affecting the inflation and

money equations have a structural interpretation. One natural extension of their work would be to establish the accuracy of their models in out-of-sample forecasting, much as we do in this paper. Further, *combining* the forecasts arising from DGE-based models with those issued from reduced-form models would likely improve the forecasting ability of either type of model.[26]

It might be, however, that the dynamics implied by the money and inflation shocks in Ireland (2001a, b) and Dib (2001, 2002) do not fit observed ones very well. Were such a situation to occur, the set-up presented here, where the shocks to the variables do not necessarily have to be given a structural interpretation that constrains the estimation, would suggest a second area of future research. One could add to the structural model the empirical representation of money and inflation dynamics favoured by a central bank. As an example, the view of money as a reliable indicator of future price pressures (embodied in the M1 vector-error-correction model (M1-VECM) used by the Bank[27]) could be added and the resulting hybrid model could become an important building block of future policy models.

Another area for future work concerns the assumption of trend stationarity that is central to the solution of the hybrid model. Voluminous empirical literature has shown that output and its components can be well characterized by I(1) processes. King, Plosser, and Rebelo (1988b) show how to modify their solution algorithms to accommodate this view within a calibrated RBC environment. Ireland (2001b) compares the out-of-sample forecasting ability of a hybrid model based on trend-stationary technology shocks with that of another based on difference stationary shocks. He presents evidence that the trend-stationary environment is the more powerful one. Future work could thus present an alternative comparison to the environment described here, where a hybrid model based on difference-stationary technology shocks would be compared to a benchmark VAR also specified in first differences and possibly including cointegration. Some authors have argued that output is also well characterized by a trend-stationary process having experienced a few breaks in its trend (Perron 1989). It would be interesting to extend the analysis to accommodate that view.

6. Conclusion

This paper has reported evidence that a hybrid model, in which a simple RBC model is augmented with a non-structural vector-process of residuals, identifies stable parameters of the RBC structure and has better forecasting accuracy than a simple,

[26] Clements and Hendry (1998, chapter 10) discuss the idea that combining the forecasts arising from two alternative models might lead to a better performance than using either model alone. They also present simple combination techniques. Li and Tkacz (2001) analyze more complex methods of combining alternative forecasts.

[27] See Adam and Hendry (2000) for a description of the M1-VECM.

reduced-form VAR.

This evidence suggests that DGE models can be framed in an econometric specification that results in a strong fit with the data and good out-of-sample forecasting ability. When coupled with their natural strength in structuring a macroeconomic model, this conformity with the data suggests that hybrid models could be important building blocks for future policy models at central banks.

References

Adam, C. and S. Hendry. 2000. "The M1 Vector-Error-Correction Model: Some Extensions and Applications." In *Money, Monetary Policy, and Transmission Mechanisms*, 151-80. Proceedings of a conference held by the Bank of Canada, November 1999. Ottawa: Bank of Canada.

Altig, D.E., C.T. Carlstrom, and K.J. Lansing. 1995. "Computable General Equilibrium Models and Monetary Policy Advice." *Journal of Money, Credit and Banking* 27(4): 1472-93.

Black, R., D. Laxton, D. Rose, and R. Tetlow. 1994. *The Bank of Canada's New Quarterly Projection Model Part 1; The Steady-State Model: SSQPM.* Technical Report No. 72. Ottawa: Bank of Canada.

Blanchard, O.J. and C.M. Kahn. 1980. "The Solution of Linear Difference Models Under Rational Expectations." *Econometrica* 48(5): 1305-12.

Blanchard, O.J. and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *The American Economic Review* 79(4): 655-73.

Chan, T. and R. Lafrance. 2001. "A Day at the Races: Forecasting the Exchange Rate." Bank of Canada. Unpublished Manuscript.

Chong, Y. and D. Hendry. 1986. "Econometric Evaluation of Linear Macro-Economic Models." *Review of Economic Studies* 53: 671-90.

Chow, G.C. and Y.K. Kwan. 1998. "How the Basic RBC Model Fails to Explain US Time Series." *Journal of Monetary Economics* 41 (2): 301-18.

Clements, M.P. and D.F. Hendry. 1998. *Forecasting Economic Time Series.* Cambridge: Cambridge University Press.

——. 1999. *Forecasting Non-stationary Economic Time Series.* Cambridge: The MIT Press.

Cogley, P. and J.M. Nason. 1995. "Output Dynamics in Real-Business-Cycle Models." *American Economic Review* 85(3): 492-511.

Coletti, D., B. Hunt, D. Rose, and R. Tetlow. 1996. *The Bank of Canada's New*

Quarterly Projection Model Part 3; The Dynamic Model: QPM. Technical Report No. 75. Ottawa: Bank of Canada.

DeJong, D.N., B.F. Ingram, and C.H. Whiteman. 2000. "A Bayesian Approach to Dynamic Macroeconomics." *Journal of Econometrics* 98: 203-23.

Dib, A. 2001. "An Estimated Canadian DSGE Model with Nominal and Real Rigidities." Bank of Canada Working Paper No. 2001-26.

——. 2002. "Nominal Rigidities and Monetary Policy in Canada Since 1981." Bank of Canada. Forthcoming working paper.

Diebold, F.X. and R.S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13(3): 253-63.

Hamilton, J.D. 1994. *Time Series Analysis.* New Jersey: Princeton University Press.

Ingram, B.F. and C.H. Whiteman. 1994. "Supplanting the 'Minnesota' Prior: Forecasting Macroeconomic Time-Series Using Real Business Cycle Models Priors." *Journal of Monetary Economics* 34: 497-510.

Ireland, P.N. 1995. "Using the Permanent Income Hypothesis for Forecasting." Federal Reserve Bank of Richmond *Economic Quarterly* 81-1 (Winter): 49-63.

——. 1997. "A Small Structural, Quarterly Model for Monetary Policy Evaluation." *Carnegie-Rochester Conference Series on Public Policy* 47: 83-108.

——. 1999. "A Method for Taking Models to the Data." Working Paper No. 9903, Federal Reserve Bank of Cleveland.

——. 2001a. "Sticky-Price Models of the Business Cycle: Specification and Stability." *Journal of Monetary Economics* 47: 3-18.

——. 2001b. "Technology Shocks and the Business Cycle: An Empirical Investigation." *Journal of Economic Dynamics and Control* 25: 703-19.

Kim, J. 2000. "Constructing and Estimating a Realistic Optimizing Model of Monetary Policy." *Journal of Monetary Economics* 45(2): 329-59.

King, R.G., C.I. Plosser, and S.T. Rebelo. 1988a. "Production, Growth and Business Cycles. I. The Basic Neoclassical Model." *Journal of Monetary Economics* 21: 195-232.

——. 1988b. "Production, Growth and Business Cycles. II. New Directions." *Journal of Monetary Economics* 21: 309-41.

King, R.G. and M.W. Watson. 1998. "The Solution of Linear Difference Systems under Rational Expectations." *International Economic Review* 39(4): 1015-26.

Leeper, E.M. and C.A. Sims. 1994. "Toward a Modern Macroeconomic Model Usable for Policy Analysis." *NBER Macroeconomics Annual* 9: 81-139.

Li, F., and G. Tkacz. 2001. "Evaluating Linear and Non-Linear Time-Varying Forecast-Combination Methods." Bank of Canada Working Paper No. 2001-12.

McGrattan, E., R. Rogerson, and R. Wright. 1997. "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy." *International Economic Review* 33: 573-601.

Moran, K. 2002. "Notes on Solving DGE Models with Linear Approximation Techniques." Bank of Canada Unpublished Manuscript.

Perron, P. 1989. "The Great Crash, the Oil Price Shock and the Unit Root Hypothesis." *Econometrica* 57: 1361-1401.

Prescott, E.C. 1986. "Theory Ahead of Business-Cycle Measurement." *Carnegie-Rochester Conference Series on Public Policy* 25: 11-44.

Sargent, T.J. 1989. "Two Models of Measurements and the Investment Accelerator." *Journal of Political Economy* 97: 251-87.

Watson, M.W. 1993. "Measures of Fit for Calibrated Models." *Journal of Political Economy* 101: 1011-41.

Table 1: Full-Sample Estimates, Standard Errors, and t-statistics

Parameter	Estimate	Standard error	t-statistic
Structural parameters			
ψ	0.0031	0.000045	69.6520
α	0.2914	0.0118	24.7072
η	1.0038	0.0005	1854.1325
A	9.4864	0.6862	13.8239
ρ	0.9860	0.0499	19.7398
σ_ξ	0.0054	0.0018	3.0323
Non-structural parameters			
d_{yy}	0.9392	0.1293	7.2653
d_{yc}	0.4930	0.2244	2.1964
d_{yh}	-0.4104	0.2561	-1.6021
d_{cy}	0.0914	0.0632	1.4456
d_{cc}	0.8660	0.0892	9.7135
d_{ch}	-0.0949	0.0932	-1.0184
d_{hy}	-0.0249	0.1884	-0.1320
d_{hc}	0.5288	0.0997	5.3063
d_{hh}	0.5935	0.1017	5.8369
v_y	0.0089	0.0017	5.1635
v_c	0.0076	0.0013	5.7378
v_h	0.0024	0.0039	0.6178
v_{yc}	0.000033	0.000023	1.4824
v_{yh}	-0.000019	0.000030	-0.6474
v_{ch}	-0.000002	0.000006	-0.2727

Table 2: Forecast Variance Decomposition (Percentage of Variance caused by Technology Shocks)

Quarters ahead	Output	Consumption	Investment	Hours worked
1	51.9213	19.9500	38.5487	83.5077
4	36.4458	27.2651	31.0333	47.1584
8	32.3095	35.2832	25.1444	31.0766
12	32.3289	41.8723	23.0418	25.0008
40	44.2223	68.1502	23.3799	18.4085
∞	53.2078	79.0914	24.5144	18.7312

Table 3: Two Subsamples: Estimates and Standard Errors

Parameter	Estimate	Standard error	Estimate	Standard error
	Subsample 1; 64Q1-82Q3		Subsample 2; 82Q4-2001Q1	
Structural parameters				
ψ	0.0032	0.000040	0.0032	0.000092
α	0.2723	0.0111	0.2909	0.0195
η	1.0042	0.0008	1.0029	0.0002
A	10.6643	0.8492	10.3890	1.3290
ρ	0.9874	0.0238	0.9265	0.0446
σ_ξ	0.0059	0.0016	0.0025	0.0005
Non-structural parameters				
d_{yy}	0.7086	0.1368	0.9884	0.1074
d_{yc}	0.7431	0.2102	0.5694	0.3214
d_{yh}	-0.1350	0.1454	-0.5595	0.2076
d_{cy}	0.0305	0.0667	0.1790	0.0586
d_{cc}	0.8023	0.1629	0.7148	0.2065
d_{ch}	-0.0669	0.0748	-0.0610	0.1233
d_{hy}	-0.1239	0.1024	-0.0228	0.0746
d_{hc}	0.5746	0.1625	0.6038	0.1256
d_{hh}	0.8221	0.1151	0.4423	0.1030
v_y	0.0090	0.0026	0.0095	0.0011
v_c	0.0091	0.0013	0.0061	0.0005
v_h	0.0030	0.0030	0.0025	0.0011
v_{yc}	0.000052	0.000026	0.000031	0.000007
v_{yh}	-0.000023	0.000024	-0.000009	0.000006
v_{ch}	-0.000005	0.000011	0.000008	0.0000004

Table 4: Two Subsamples: Tests for Parameter Stability

Stability of all 21 estimated parameters: 95% confidence interval: 10.3-35.5 99% confidence interval: 8.0-41.4	Test statistic 55.4071
Stability of the 6 structural parameters: 95% confidence interval: 1.24-14.4 99% confidence interval: 0.7-18.5	13.7242
Stability of the remaining 15 parameters: 95% confidence interval: 6.26-27.5 99% confidence interval: 4.6-32.8	32.0013

Table 5: Forecasting Accuracy: Comparison of the Structural Model with a Two-Lag, Reduced-Form VAR, 1985Q1-2001Q1

Quarters ahead	1	2	3	4
Panel A: Output				
RMSE hybrid model	$1.11622^{D,F}$	$1.80975^{D,F}$	$2.47052^{D,F}$	$3.06950^{D,F}$
RMSE reduced-form model	1.22660	2.24965	3.31399	4.25331
Panel B: Consumption				
RMSE hybrid model	0.67731^{F}	0.95901^{D}	1.23302^{D}	1.58780^{D}
RMSE reduced-form model	0.71490	1.15446	1.63320	2.16990
Panel C: Investment				
RMSE hybrid model	4.33331^{F}	6.34281^{F}	$8.29900^{D,F}$	$9.87157^{D,F}$
RMSE reduced-form model	4.50063	7.17825	10.07824	12.41626
Panel D: Hours worked				
RMSE hybrid model	$0.52833^{D,F}$	$0.82621^{D,F}$	1.15416^{D}	1.43775^{D}
RMSE reduced-form model	0.58010	1.02618	1.61342	2.17836

Notes:

Superscript D signals that the model performs better at the 5 per cent significance level according to the Diebold and Mariano test.

Superscript F signals that the model performs better at the 5 per cent significance level according to the forecast-encompassing test.

Table 6: Forecasting Accuracy: Comparison of the Structural Model with a Two-Lag, Reduced-Form VAR, 1990Q1-2001Q1

Quarters ahead	1	2	3	4
Panel A: Output				
RMSE hybrid model	$1.14328^{D,F}$	$1.86682^{D,F}$	$2.56218^{D,F}$	$3.15684^{D,F}$
RMSE reduced-form model	1.31746	2.35840	3.39280	4.23789
Panel B: Consumption				
RMSE hybrid model	0.71129	0.99388^{F}	1.32742^{F}	1.70904^{F}
RMSE reduced-form model	0.64525	0.93417	1.30660	1.74553
Panel C: Investment				
RMSE hybrid model	4.52747^{F}	$6.63148^{D,F}$	$8.68483^{D,F}$	$10.25775^{D,F}$
RMSE reduced-form model	4.99117	8.02066	11.17249	13.45428
Panel D: Hours worked				
RMSE hybrid model	$0.54691^{D,F}$	$0.83699^{D,F}$	1.16836^{D}	1.44166^{D}
RMSE reduced-form model	0.60755	0.99681	1.52713	1.99167

Notes:

Superscript D signals that the model performs better at the 5 per cent significance level according to the Diebold and Mariano test.

Superscript F signals that the model performs better at the 5 per cent significance level according to the forecast-encompassing test.

Table 7: Forecasting Accuracy: Comparison of the Structural Model with a One-Lag, Reduced-Form VAR, 1985Q1-2001Q1

Quarters ahead	1	2	3	4
Panel A: Output				
RMSE hybrid model	$1.11622^{D,F}$	$1.80975^{D,F}$	2.47052^{D}	3.06950^{D}
RMSE reduced-form model	1.35566	2.29463	3.15159	3.88320
Panel B: Consumption				
RMSE hybrid model	0.67731^{F}	0.95901^{F}	1.23302	1.58780
RMSE reduced-form model	0.71918	1.09978	1.46889	1.89112
Panel C: Investment				
RMSE hybrid model	$4.33331^{D,F}$	$6.34281^{D,F}$	$8.29900^{D,F}$	$9.87157^{D,F}$
RMSE reduced-form model	5.06337	7.75275	10.11258	11.96880
Panel D: Hours worked				
RMSE hybrid model	$0.52833^{D,F}$	$0.82621^{D,F}$	$1.15416^{D,F}$	1.43775^{D}
RMSE reduced-form model	0.69454	1.22180	1.75604	2.27218

Notes:

Superscript D signals that the model performs better at the 5 per cent significance level according to the Diebold and Mariano test.

Superscript F signals that the model performs better at the 5 per cent significance level according to the forecast-encompassing test.

Figure 1: In-Sample Forecasting Errors

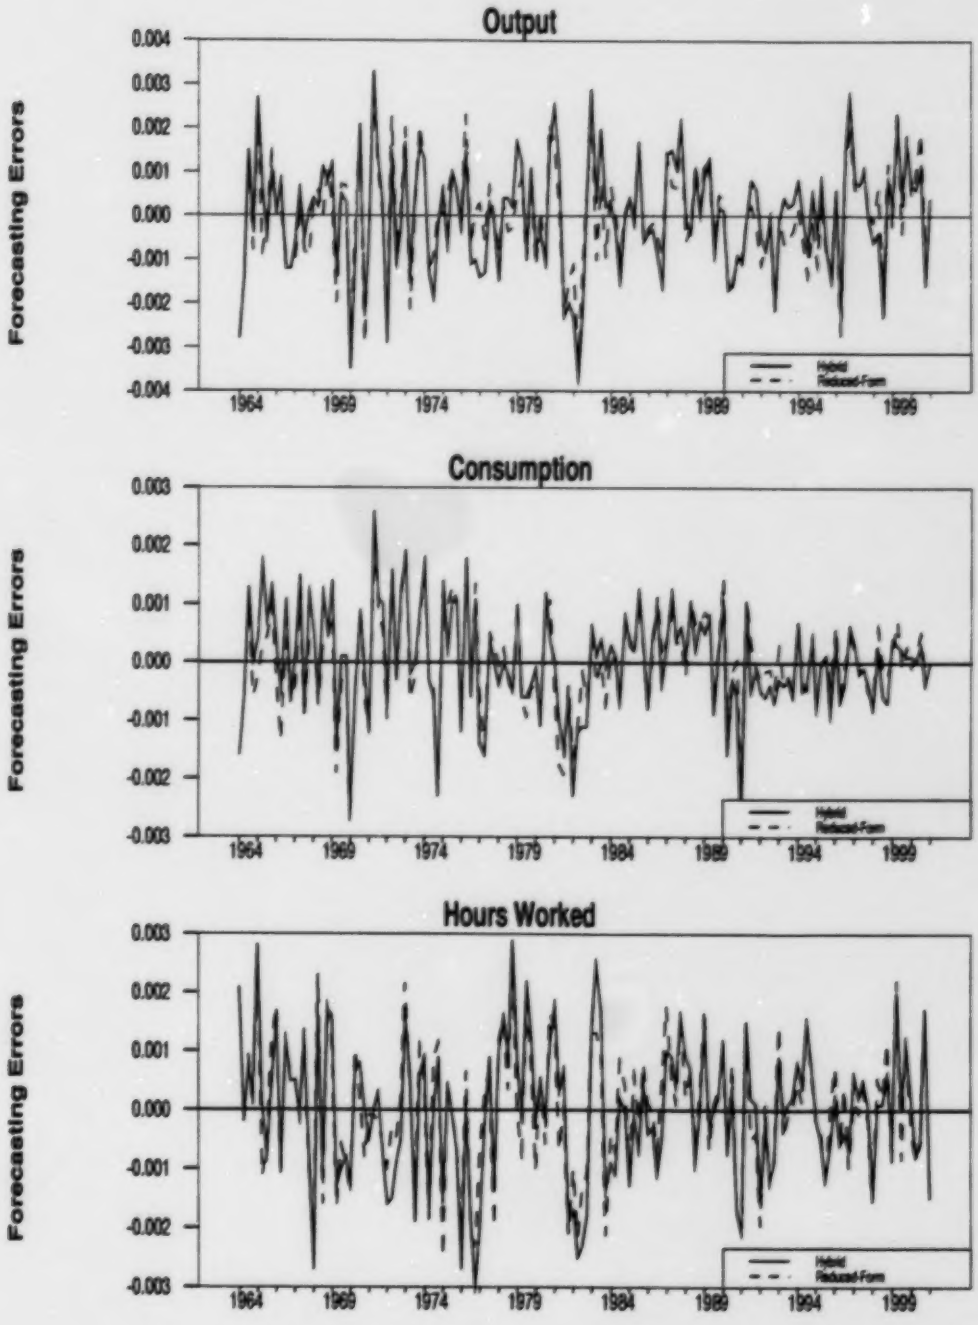

Figure 2: Out-of-Sample Forecasting Errors, 1 Quarter Ahead

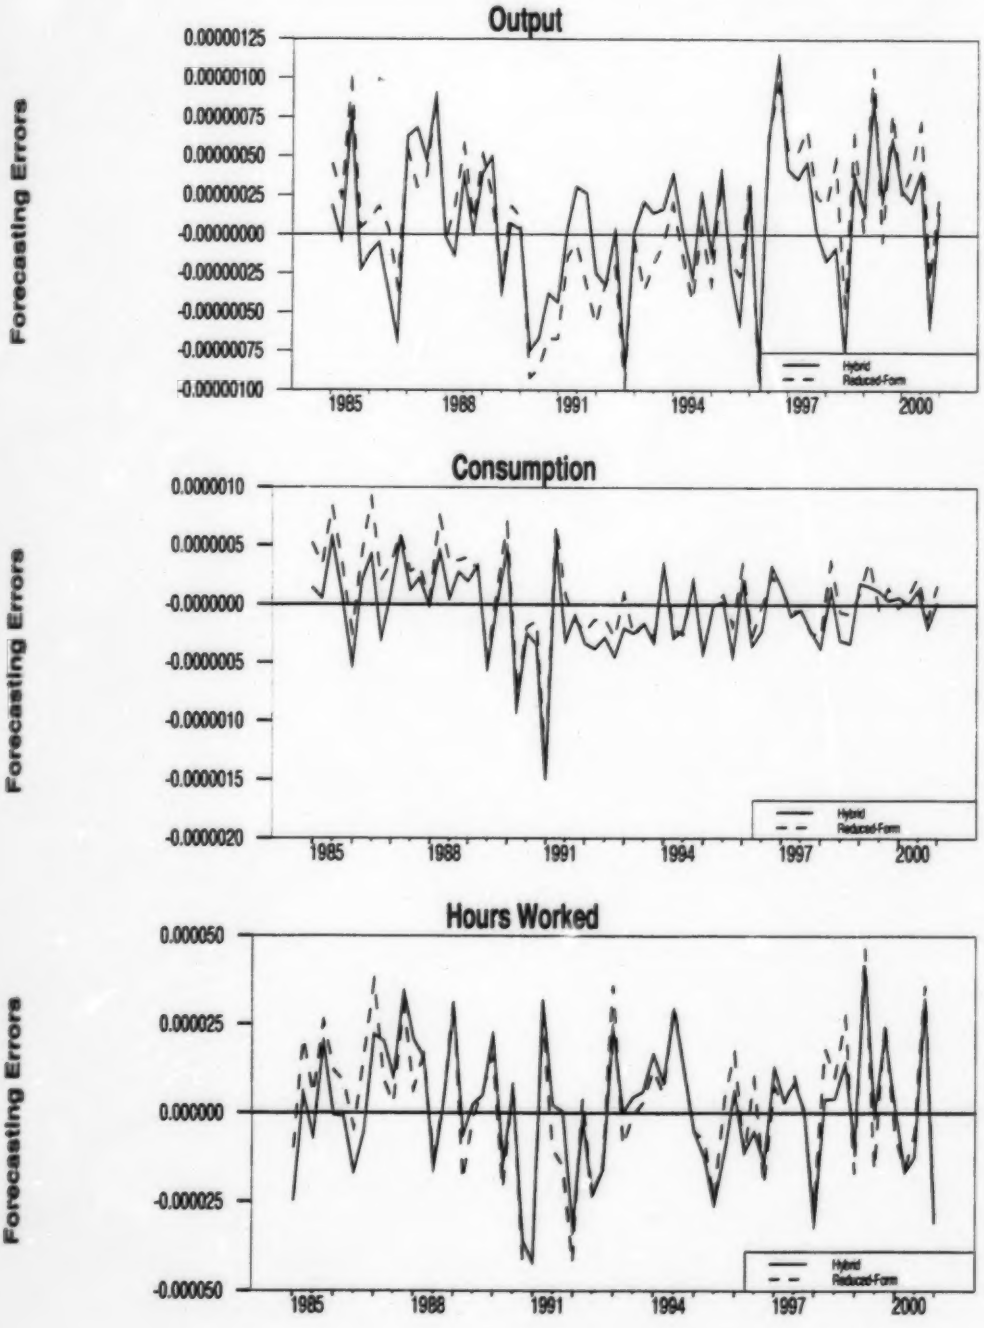

Figure 3: Out-of–Sample Forecasting Errors, 4 Quarters Ahead

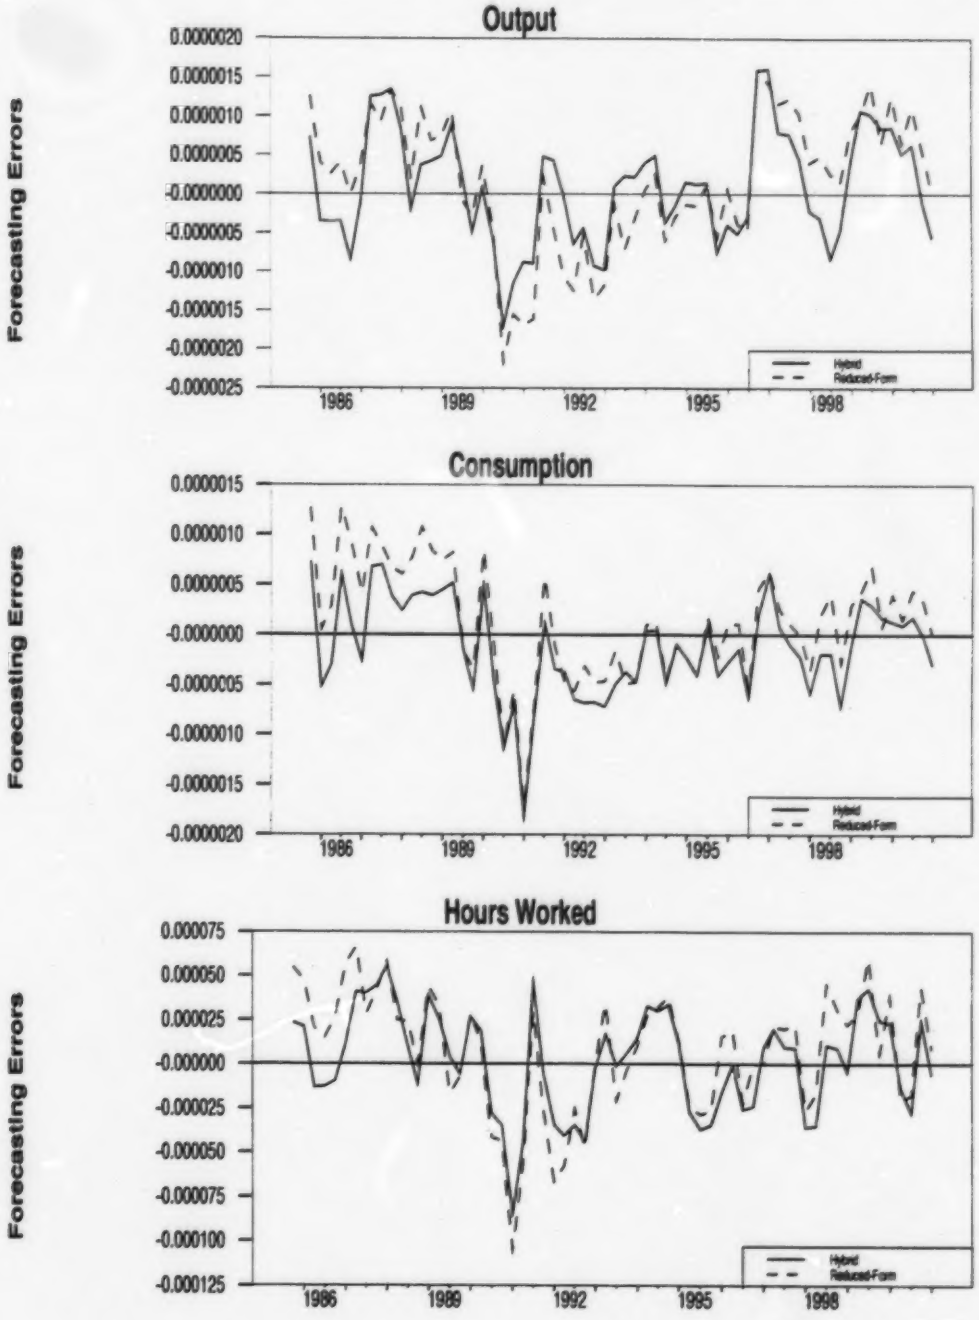

Appendix A: Computing the Log-Likelihood of the Model using the Kalman Filter: A Brief Overview

Recall our empirical model from (18):

$$\begin{aligned}\begin{bmatrix} \mathbf{s}_{t+1} \\ \mathbf{u}_{t+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{M}_1 & 0 \\ 0 & \mathbf{D} \end{bmatrix} \begin{bmatrix} \mathbf{s}_t \\ \mathbf{u}_t \end{bmatrix} + \begin{bmatrix} \mathbf{M}_2 & 0 \\ 0 & \mathbf{I} \end{bmatrix} \begin{bmatrix} \epsilon_{t+1} \\ \nu_{t+1} \end{bmatrix} \\ \mathbf{f}_t &= \begin{bmatrix} \mathbf{\Pi} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{s}_t \\ \mathbf{u}_t \end{bmatrix}.\end{aligned} \tag{A.1}$$

Compare this system with the one described in Hamilton's (1993, chapter 13) discussion of state space models and the Kalman filter:

$$\begin{aligned} \xi_{t+1} &= \mathbf{F} \cdot \xi_t + \mathbf{v}_{t+1}; \\ \mathbf{y}_t &= \mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \xi_t + \mathbf{w}_t; \\ E(\mathbf{v}_t \mathbf{v}_t') &= \mathbf{Q}; \\ E(\mathbf{w}_t \mathbf{w}_t') &= \mathbf{R}. \end{aligned}$$

The equivalence between the two systems is established by defining $\mathbf{y}_t = \mathbf{f}_t$, $\mathbf{x}_t = 0$, $\xi_t = [\mathbf{s}_t' \ \mathbf{u}_t']'$, $\mathbf{w}_t = 0$, $\mathbf{v}_t = [\mathbf{M}_2 \epsilon_t' \ \nu_t']'$ as well as the following matrices:

$$\mathbf{A} = 0; \ \mathbf{H} = \begin{bmatrix} \mathbf{\Pi} & \mathbf{I} \end{bmatrix}; \ \mathbf{F} = \begin{bmatrix} \mathbf{M}_1 & 0 \\ 0 & \mathbf{D} \end{bmatrix}; \ \mathbf{Q} = \begin{bmatrix} \sigma_\epsilon^2 \mathbf{M}_2' \mathbf{M}_2 & 0 \\ 0 & \mathbf{V} \end{bmatrix}; \ \mathbf{R} = 0.$$

The Kalman filter is used to compute the best forecast of the unobserved state, ξ_t, conditional on information available at time *t-1*. Denote this forecast by $\widehat{\xi}_{t|t-1}$. Further, denote the MSE of this forecast by $\mathbf{P}_{t|t-1}$.[28] Conditional on starting values $\hat{\xi}_{1|0}$ and $P_{1|0}$, the following recursive structure describing the evolution of $\hat{\xi}_{t+1|t}$ and $P_{t+1|t}$ emerges:

$$\mathbf{K}_t = \mathbf{F}\mathbf{P}_{t|t-1}\mathbf{H}(\mathbf{H}'\mathbf{P}_{t|t-1}\mathbf{H})^{-1}; \tag{A.2}$$

$$\widehat{\xi}_{t+1|t} = \mathbf{F}\hat{\xi}_{t|t-1} + \mathbf{K}_t(y_t - \mathbf{H}'\hat{\xi}_{t|t-1}); \tag{A.3}$$

$$\mathbf{P}_{t+1|t} = (\mathbf{F} - \mathbf{K}_t\mathbf{H}')\mathbf{P}_{t|t-1}(\mathbf{F}' - \mathbf{H}\mathbf{K}_t') + \mathbf{Q}. \tag{A.4}$$

The intuition behind this updating sequence is that, at each step, the econometrician uses the *observed* forecasting errors $(y_t - \mathbf{H}' \cdot \widehat{\xi}_{t|t-1})$ and knowledge of the parametric form of the system to update the best estimate of the unobserved states, ξ_t. The mechanics of this updating takes the form of linear projection and is detailed in Hamilton.

[28] So that $\mathbf{P}_{t|t-1} = E_{t-1}[(\xi_t - \widehat{\xi}_{t|t-1})(\xi_t - \widehat{\xi}_{t|t-1})']$.

Under the assumption that the errors are normally distributed, the conditional likelihood of $\mathbf{y}_t$ is normal, and given by:

$$f(\mathbf{y}_t|I_{t-1},\ \theta) = (2\pi)^{-0.5n}|\mathbf{H}'\mathbf{P}_t|_{t-1}\mathbf{H}|^{-0.5}exp\left[(\mathbf{y}_t - \mathbf{H}'\cdot\widehat{\xi}_t|_{t-1})'(\mathbf{H}'\mathbf{P}_t|_{t-1}\mathbf{H})^{-1}(\mathbf{y}_t - \mathbf{H}'\cdot\widehat{\xi}_t|_{t-1})\right], \tag{A.5}$$

where the likelihood is indexed by θ, the vector of parameters, to remind the reader that the matrices $\mathbf{F}, \mathbf{H}, \mathbf{Q}, \mathbf{K}_t$, and $\mathbf{P}_t|_{t-1}$ are all functions of this vector. Summing this sequence of conditional likelihoods gives the log-likelihood for the complete sample $\mathbf{Y}_t$:

$$logL(\mathbf{Y}_t|\theta) = \sum_{t=0}^{T} logf(\mathbf{y}_t|I_{t-1},\ \theta). \tag{A.6}$$

This expression is maximized with respect to θ to deliver the maximum-likelihood estimate of the model. Note that we impose the presence of non-explosive roots in both matrices $\mathbf{M}_1$ and $\mathbf{\Pi}$ by assigning a very low value to the likelihood when the algorithm tries values of the parameter vector that imply such explosive roots.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.***

2002

2002-17	Does Exchange Rate Policy Matter for Growth?	J. Bailliu, R. Lafrance, and J.-F. Perrault
2002-16	A Market Microstructure Analysis of Foreign Exchange Intervention in Canada	C. D'Souza
2002-15	Corporate Bond Spreads and the Business Cycle	Z. Zhang
2002-14	Entrepreneurship, Inequality, and Taxation	C.A. Meh
2002-13	Towards a More Complete Debt Strategy Simulation Framework	D.J. Bolder
2002-12	Modelling Financial Instability: A Survey of the Literature	A. Lai
2002-11	Risk, Entropy, and the Transformation of Distributions	R.M. Reesor and D.L. McLeish
2002-10	La fiabilité des estimations de l'écart de production au Canada	J.-P. Cayen et S. van Norden
2002-9	The Microstructure of Multiple-Dealer Equity and Government Securities Markets: How They Differ	T. Gravelle
2002-8	Restructuring in the Canadian Economy: A Survey of Firms	C. Kwan
2002-7	Contribution of ICT Use to Output and Labour-Productivity Growth in Canada	H. Khan and M. Santos
2002-6	Currency Fluctuations, Liability Dollarization, and the Choice of Exchange Rate Regimes in Emerging Markets	P.N. Osakwe
2002-5	The Effects of Bank Consolidation on Risk Capital Allocation and Market Liquidity	C. D'Souza and A. Lai
2002-4	Does Micro Evidence Support the Wage Phillips Curve in Canada?	J. Farès
2002-3	An Introduction to Wavelets for Economists	C. Schleicher
2002-2	Asset Allocation Using Extreme Value Theory	Y. Bensalah
2002-1	Taylor Rules in the Quarterly Projection Model	J. Armour, B. Fung, and D. Maclean

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca